

	Type	L #	Hits	Search Text	DBs	Time Stamp	Comments	Error or Definition	Errors
1	IS&R	L1	3100	(710/10,3,100,240).CCLS	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 09:55			
2	BRS	L2	0	1 and "network node" and bus and "programmable switch"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 09:56			
3	BRS	L3	61	1 and "network node" and bus	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 09:56			
4	BRS	L4	0	3 and programmable near3 switch	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:12			
5	BRS	L5	1963917	3 switch	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 09:57			
6	BRS	L6	0	3 and programmable near3 switch	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 09:57			
7	BRS	L7	32	3 and switch	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 09:58			
8	BRS	L8	25535915	@ad<"20040421"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 09:58			
9	BRS	L10	15042	process adj block	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 09:59			

10	BRS	L11	18278	processing adj block	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 09:59			
----	-----	-----	-------	----------------------	---	---------------------	--	--	--

	Type	L #	Hits	Search Text	DBs	Time Stamp	Comments	Error or Definitio	Errors
11	BRS	L12	2	9and 11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 09:59			
12	BRS	L13	0	9 and 11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:01			
13	BRS	L14	736	spare near3 processing	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:00			
14	BRS	L15	17	14 and processing near3 defective	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:01			
15	BRS	L16	26	9 and "17"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:01			
16	BRS	L17	0	9 and 15	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:02			
17	BRS	L18	12	8 and 15	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:02			
18	BRS	L9	32	7 and 8	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:04			
19	BRS	L19	6863	programmable near3 switch	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:12			

20	BRS	L20	6863	programmable near3 switch	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:12			
----	-----	-----	------	------------------------------	---	---------------------	--	--	--

	Type	L #	Hits	Search Text	DBs	Time Stamp	Comments	Error or Definition	Errors
21	BRS	L21	2900	20 and network	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:12			
22	BRS	L22	2036	21 and processing	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:12			
23	BRS	L23	386	22 and input adj port	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:13			
24	BRS	L24	333	23 and output adj port	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:13			
25	BRS	L25	307	24 and 8	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:13			
26	BRS	L26	14	25 and connection adj information	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/05/03 10:14			



[Subscribe](#) (Full Service) [Register](#) (Limited Service, Free) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

"processing block" and network and node and port and memor

SEARCH



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used **processing**

block and **network** and **node** and **port** and **memory** and **bus** and **programmable switch** and

Found 37,480 of 175,083

Sort results by

relevance



[Save results to a Binder](#)

[Try an Advanced Search](#)

Display results

expanded form



[Search Tips](#)

[Try this search in The ACM Guide](#)

☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Balancing performance and flexibility with hardware support for network architectures](#)



Ilija Hadžić, Jonathan M. Smith

November 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 4

Publisher: ACM Press

Full text available: pdf(719.03 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The goals of performance and flexibility are often at odds in the design of network systems. The tension is common enough to justify an architectural solution, rather than a set of context-specific solutions. The Programmable Protocol Processing Pipeline (P4) design uses programmable hardware to selectively accelerate protocol processing functions. A set of field-programmable gate arrays (FPGAs) and an associated library of network processing modules implemented in hardware are augmented with so ...

Keywords: FPGA, P4, computer networking, flexibility, hardware, performance, programmable logic devices, programmable networks, protocol processing

2 [Compiler-directed channel allocation for saving power in on-chip networks](#)



Guangyu Chen, Feihui Li, Mahmut Kandemir

January 2006 **ACM SIGPLAN Notices , Conference record of the 33rd ACM SIGPLAN-SIGACT symposium on Principles of programming languages POPL '06**, Volume 41 Issue 1

Publisher: ACM Press

Full text available: pdf(943.11 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Increasing complexity in the communication patterns of embedded applications parallelized over multiple processing units makes it difficult to continue using the traditional bus-based on-chip communication techniques. The main contribution of this paper is to demonstrate the importance of compiler technology in reducing power consumption of applications designed for emerging multi processor, NoC (Network-on-Chip) based embedded systems. Specifically, we propose and evaluate a compiler-directed a ...

Keywords: NoC, compiler, energy consumption



The V distributed system

David Cheriton

March 1988 **Communications of the ACM**, Volume 31 Issue 3

Publisher: ACM Press

Full text available: [pdf\(2.55 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The V distributed System was developed at Stanford University as part of a research project to explore issues in distributed systems. Aspects of the design suggest important directions for the design of future operating systems and communication systems.



4 Reprogrammable network packet processing on the field programmable port extender (FPX)



John W. Lockwood, Naji Naufel, Jon S. Turner, David E. Taylor

February 2001 **Proceedings of the 2001 ACM/SIGDA ninth international symposium on Field programmable gate arrays**

Publisher: ACM Press

Full text available: [pdf\(257.98 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A prototype platform has been developed that allows processing of packets at the edge of a multi-gigabit-per-second network switch. This system, the Field Programmable Port Extender (FPX), enables packet processing functions to be implemented as modular components in reprogrammable hardware. All logic on the on the FPX is implemented in two Field Programmable Gate Arrays (FPGAs). Packet processing functions in the system are implemented as dynamically-loadable modules. Core functi ...

Keywords: ATM, FPGA, IP, Internet, hardware, modularity, network, packet, processing, reconfiguration, routing

5 YARDS: FPGA/MPU hybrid architecture for telecommunication data processing



Akihiro Tsutsui, Toshiaki Miyazaki

February 1997 **Proceedings of the 1997 ACM fifth international symposium on Field-programmable gate arrays**

Publisher: ACM Press

Full text available: [pdf\(1.04 MB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

6 Monitoring and performance measuring distributed systems during operation



D. Wybraniec, D. Haban

May 1988 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1988 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '88**, Volume 16 Issue 1

Publisher: ACM Press

Full text available: [pdf\(1.22 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes an integrated tool for monitoring distributed systems continuously during operation. A hybrid monitoring approach is used. As special hardware support a test and measurement processor (TMP) was designed, which is part of each node in an experimental multicomputer system. Each TMP runs local parts of the monitoring software for its node, while all the TMPs are connected to a central test station via a separate TMP interconnection network. The monitoring system is transpa ...

7

A message passing coprocessor for distributed memory multicomputers



Jiun-Ming Hsu, Prithviraj Banerjee

November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society

Full text available:  [pdf\(1.25 MB\)](#) · Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

This paper presents the architecture, methodology and performance evaluation of a *message passing coprocessor* (MPC) which can accelerate message communication in a distributed memory multicomputer. The MPC is a microprogrammable processor which off-loads the CPU of the burden of communication and speeds up the software processing by directly executing message passing instructions in microcode. It supports process scheduling, message buffer management, and fast buffer copying. The most uni ...

8 Run-time adaptation in river



Remzi H. Arpaci-Dusseau

February 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(849.04 KB\)](#) · Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present the design, implementation, and evaluation of run-time adaptation within the River dataflow programming environment. The goal of the River system is to provide adaptive mechanisms that allow database query-processing applications to cope with performance variations that are common in cluster platforms. We describe the system and its basic mechanisms, and carefully evaluate those mechanisms and their effectiveness. In our analysis, we answer four previously unanswered and important que ...

Keywords: Performance availability, clusters, parallel I/O, performance faults, robust performance, run-time adaptation

9 Distributed operating systems



Andrew S. Tanenbaum, Robbert Van Renesse

December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(5.49 MB\)](#) · Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...

10 Special session on reconfigurable computing: Adaptive architectures for an OTN processor: reducing design costs through reconfigurability and multiprocessing



Tudor Murgan, Mihail Petrov, Mateusz Majer, Peter Zipf, Manfred Glesner, Ulrich Heinkel, Joerg Pleickhardt, Bernd Bleisteiner

April 2004 **Proceedings of the 1st conference on Computing frontiers**

Publisher: ACM Press

Full text available:  [pdf\(1.01 MB\)](#) · Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The standardisation process of Optical Transport Networks generally spans a long period of time. For providers intending to be present early on the market, this implies costly design re-spins if the wrong "flavour" of the protocol standard has been implemented. Extending a protocol processing device through application specific reconfigurable elements or multiprocessor units augment its flexibility. Thus, the architecture can be

upgraded to standard updates or changes not even considered at desi ...

Keywords: ITU-T G.709, multiprocessor and reconfigurable architectures, optical transport networks, standard upgrades

11 Novel FPGA applications: CUSP: a modular framework for high speed network



applications on FPGAs

Graham Schelle, Dirk Grunwald

February 2005 **Proceedings of the 2005 ACM/SIGDA 13th international symposium on Field-programmable gate arrays**

Publisher: ACM Press

Full text available: pdf(547.03 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

For several years now, modern FPGAs have included onchip network related hard cores. These cores include Xilinx's RocketIO and Altera's RapidIO serial transceivers. However, to use these cores in a complete networking application may be a daunting task to a non-networking expert. In addition to the complicated use of these components, the high performance needs of modern networking applications require designs that are optimized for low latency and a moderately high clock rate. Therefore to meet ...

Keywords: networking, parallelism, reconfigurable hardware, speculation

12 The design of a distributed kernel



David R. Cheriton

January 1981 **Proceedings of the ACM '81 conference**

Publisher: ACM Press

Full text available: pdf(668.86 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The design of a distributed kernel for a multi-processor machine is described that combines the advantages of a shared centralized kernel with the efficiency of separate kernels per processor. The base machine architecture is a star network of microcomputer modules with a minicomputer as the central node, implemented using off-the-shelf hardware. The kernel implements a uniform, location transparent model of processes communicating via messages. Preliminary measurements are given for the me ...

13 Evaluation of the Raw Microprocessor: An Exposed-Wire-Delay Architecture for ILP and Streams



Michael Bedford Taylor, Walter Lee, Jason Miller, David Wentzlaff, Ian Bratt, Ben Greenwald, Henry Hoffmann, Paul Johnson, Jason Kim, James Psota, Arvind Saraf, Nathan Shnidman, Volker Strumpfen, Matt Frank, Saman Amarasinghe, Anant Agarwal

March 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture ISCA '04,**
Volume 32 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(376.05 KB) Additional Information: [full citation](#), [abstract](#), [citations](#)

This paper evaluates the Raw microprocessor. Raw addresses the challenge of building a general-purpose architecture that performs well on a larger class of stream and embedded computing applications than existing microprocessors, while still running existing ILP-based sequential programs with reasonable performance in the face of increasing wire delays. Raw approaches this challenge by implementing plenty of on-chip resources - including logic, wires, and pins - in a tiled arrangement, and exposing the ...

14



Programmable/extensible networks: Addressing data compatibility on programmable



network platforms

Ada Gavrilovska, Karsten Schwan
October 2005 **Proceedings of the 2005 symposium on Architecture for networking and communications systems ANCS '05**

Publisher: ACM Press

Full text available: pdf(268.88 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Large-scale applications require the efficient exchange of data across their distributed components, including data from heterogeneous sources and to widely varying clients. Inherent to such data exchanges are (1) discrepancies among the data representations used by sources, clients, or intermediate application components (e.g., due to natural mismatches or due to dynamic component evolution), and (2) requirements to route, combine, or otherwise manipulate data as it is being transferred. As a r ...

Keywords: data morphing, network processors, streaming applications

15 Efficient Field Processing Cores in an Innovative Protocol Processo System-on-Chip

G. Lykakis, N. Mouratidis, K. Vlachos, N. Nikolaou, S. Perissakis, G. Sourdis, G. Konstantoulakis, D. Pnevmatikatos, D. Reisis

March 2003 **Proceedings of the conference on Design, Automation and Test in Europe: Designers' Forum - Volume 2 DATE '03**

Publisher: IEEE Computer Society

Full text available: pdf(179.42 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)



[Publisher Site](#)

We present an innovative protocol processor component that combines wire-speed processing for low-level, and best effort processing for higher-level protocols. The component is a System-on-Chip that integrates variable size packet buffering, specialised cores for header and field processing, generic RISC cores and scheduling blocks. We focus on the main innovation, the reprogrammable pipeline module, and discuss its internal architecture, optimised to perform field processing on byte streams, as ...

16 The K2 distributed memory parallel processor: architecture, compiler, and operating



system

M. Annaratone, M. Fillo, M. Halbherr, R. Rühl, P. Steiner, M. Viredaz

August 1991 **Proceedings of the 1991 ACM/IEEE conference on Supercomputing**

Publisher: ACM Press

Full text available: pdf(1.13 MB) Additional Information: [full citation](#), [references](#), [citing](#), [index terms](#)

17 Keynote: Powering networks on chips: energy-efficient and reliable interconnect



design for SoCs

Luca Benini, Giovanni De Micheli

September 2001 **Proceedings of the 14th international symposium on Systems synthesis**

Publisher: ACM Press

Full text available: pdf(143.96 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#)

We consider *systems on chips* (SoCs) that will be designed and produced in five to ten years from today, with gate lengths in the range 50-100nm. We address the distinguishing features of a design methodology that aims at achieving reliable designs under the limitations of the interconnect technology. Specifically, we consider energy consumption reduction, under guaranteed *quality of service* (QoS), as a main objective in system design.

Keywords: low-energy design, networks, systems on chips

18 A model for recentralization of computing: (distributed processing comes home)



Harold Lorin

March 1990 **ACM SIGARCH Computer Architecture News**, Volume 18 Issue 1

Publisher: ACM Press

Full text available: [pdf\(1.38 MB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Distributed systems commonly contain heterogeneity at all levels of systems structure, differentiated by function (special servers), operating systems and architecture within a single system. On the other hand, large mainframes tend to be more homogeneous in their structures, even when they are multiprocessors. This paper explores a way of using the models of heterogeneous distributed computing within a mainframe. The argument is that appropriate restructuring of the mainframe can achieve a conv ...

19 The Vector-Thread Architecture



Ronny Krashinsky, Christopher Batten, Mark Hampton, Steve Gerding, Brian Pharris, Jared Casper, Krste Asanovic

March 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture ISCA '04**, Volume 32 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: [pdf\(317.13 KB\)](#) Additional Information: [full citation](#), [abstract](#)

The vector-thread (VT) architectural paradigm unifies the vector and multithreaded compute models. The VT abstraction provides the programmer with a control processor and a vector of virtual processors (VPs). The control processor can use vector-fetch commands to broadcast instructions to all the VPs or each VP can use thread-fetches to direct its own control flow. A seamless intermixing of the vector and threaded control mechanisms allows a VT architecture to flexibly and compactly encode application ...

20 Processor scheduling on multiprogrammed, distributed memory parallel computers



Sanjeev K. Setia, Mark S. Squillante, Satish K. Tripathi

June 1993 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1993 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '93**, Volume 21 Issue 1

Publisher: ACM Press

Full text available: [pdf\(1.39 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Multicomputers, consisting of many processing nodes connected through a high speed interconnection network, have become an important and common platform for a large body of scientific computations. These parallel systems have traditionally executed programs in batch mode, or have at most space-shared the processors among multiple programs using a static partitioning policy. This, however, can result in relatively low system utilization and throughput for important classes of scientific applicati ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)